

Particle Model Predictive Control: Tractable Stochastic Nonlinear Output-Feedback MPC

Martin A. Sehr & Robert R. Bitmead

*Department of Mechanical & Aerospace Engineering, University of
California, San Diego, La Jolla, CA 92093-0411, USA
(e-mail: {msehr, rbitmead}@ucsd.edu).*

Abstract: We combine conditional state density construction with an extension of the Scenario Approach for stochastic Model Predictive Control to nonlinear systems to yield a novel particle-based formulation of stochastic nonlinear output-feedback Model Predictive Control. Conditional densities given noisy measurement data are propagated via the Particle Filter as an approximation of the Bayesian Filter. This enables a particle-based representation of the conditional state density, or information state, which naturally merges with scenario generation including the current system state. This approach attempts to address the computational tractability questions of general nonlinear stochastic optimal control. The Particle Filter and the Scenario Approach are shown to be fully compatible and – based on the time- and measurement-update stages of the Particle Filter – incorporated into the optimization over future control sequences. A numerical example is presented and examined for the dependence of solution and computational burden on the sampling configurations of the densities, scenario generation and the optimization horizon.

Keywords: stochastic control, model predictive control, nonlinear control, information state, particle filtering.

1. INTRODUCTION

Model Predictive Control (MPC), in its original formulation, is a full-state feedback law (see Mayne et al. (2000); Mayne (2014); Maciejowski (2002)). This underpins two theoretical limitations of MPC: accommodation of output-feedback, and extension to include a compelling robustness theory given the state dimension is fixed. This paper addresses the first of these issues in a rather general, though practical setup.

There have been a number of approaches to output-feedback MPC, mostly hinging on the replacement of the measured true state by a state estimate, which is computed via Kalman filtering (e.g. Sehr and Bitmead (2016b); Yan and Bitmead (2005)), moving-horizon estimator (e.g. Copp and Hespanha (2014); Sui et al. (2008)), tube-based minimax estimators (e.g. Mayne et al. (2009)), etc. Apart from Copp and Hespanha (2014), these designs, often for linear systems, separate the estimator design from the control design. The control problem may be altered to accommodate the state estimation error by methods such as: constraint tightening as in Yan and Bitmead (2005), chance/probabilistic constraints as in Cannon et al. (2012) or Schwarm and Nikolaou (1999), and so forth. Likewise, for nonlinear problems, where the state estimation behavior is affected by control signal properties, the control may be modified to enhance the excitation properties of the estimator, as suggested in Chisci et al. (2001); Marafioti et al. (2014). Each of these aspects of accommodation is made in an isolated fashion.

The stochastic nonlinear output-feedback MPC algorithm presented in this paper is motivated by the structure of Stochastic Model Predictive Control (SMPC) via finite-horizon stochastic optimal control. The latter method requires propagating conditional state densities using a Bayesian Filter (BF) and solution of the Stochastic Dynamic Programming Equation (SDPE). By virtue of implementing a truly optimal finite-horizon control law in a receding horizon fashion, one can deduce a number of properties of the closed-loop dynamics, including recursive feasibility of the SMPC controller, stochastic stability and bounds characterizing closed-loop infinite-horizon performance, as discussed in Sehr and Bitmead (2016a).

Unfortunately, solving for the stochastic optimal output-feedback controller, even on the finite horizon, is computationally intractable except for special cases such as linear quadratic Gaussian MPC because of the need to solve the SDPE, which incorporates the duality of the optimal control law in its effect on state observability. While the BF, required to propagate the conditional state densities, is readily approximated using a Particle Filter (PF), open-loop solution of the SDPE results in the loss of the duality of the optimal control. While not discussed in this paper, this effect can be mitigated sub-optimally by imposing excitation requirements as in Chisci et al. (2001); Marafioti et al. (2014).

Approximately propagating the conditional state densities by means of the PF naturally invites combination with the more recent advances in Scenario Model Pre-

dictive Control (SCMPC), as discussed for instance by Calafiore and Fagiano (2013); Grammatico et al. (2016); Lee (2014); Mesbah et al. (2014); Schildbach et al. (2014). Scenario methods deal with optimization of difficult, non-convex problems in which the initial task is recast as a parametrized collection of simpler, generally convex problems. Random sampling of uncertain signals and parameters is performed and the resulting collection of deterministic problem instances is solved. The focus has been on full state feedback for systems with linear dynamics and probabilistic state constraints. The technical construction is to take a sufficient number of samples (scenarios) to provide an adequate reconstruction of future controlled state densities for design.

In contrast to solving the SDPE underlying the stochastic optimal control problem, the future controlled state densities in SCMPC are open-loop constructions. However, they present a natural fit combined with the particle-based conditional density approximations generated by the PF, where individual particles can be interpreted as scenarios from an estimation perspective. Moreover, while SCMPC is typically formulated in the linear case, the basic idea extends to the nonlinear case, albeit with the loss of many computation-saving features. In this paper, we propose and discuss this output-feedback version of SCMPC combined with the PF, which we call Particle Model Predictive Control (PMPC). Compared with the stochastic optimal output-feedback controller (computed via BF and SDPE), the PMPC controller is suboptimal in not accommodating future measurement updates and thereby losing both exact constraint violation probabilities along the horizon and the probing requirement inherent to stochastic optimal control. On the other hand, PMPC enables a generally applicable and, at least for small state dimensions, computationally tractable alternative for nonlinear stochastic output-feedback control.

The structure of the paper evolves as follows. We briefly introduce the problem setup in Section 2 and SMPC in Section 3 and proceed by introducing the PMPC control algorithm based on its individual components and parameters in Section 4. After describing the algorithm and its correspondence to SMPC, we use a scalar, but challenging and representative nonlinear example to demonstrate computational tractability and dependence of the proposed PMPC closed-loop behavior on a number of parameters in Section 5. The example features nonlinear state and measurement equations and probabilistic state constraints under significant measurement noise. Finally, we conclude with Section 6.

2. STOCHASTIC OPTIMAL CONTROL – SETUP

We consider receding horizon output-feedback control for nonlinear stochastic systems of the form

$$x_{t+1} = f(x_t, u_t, w_t), \quad x_0 \in \mathbb{R}^n, \quad (1)$$

$$y_t = h(x_t, v_t), \quad (2)$$

starting from known initial state density $\pi_{0|-1} = \text{pdf}(x_0)$. To this end, we denote the data available at time t by

$$\zeta^t \triangleq \{y_0, u_0, y_1, u_1, \dots, u_{t-1}, y_t\}, \quad \zeta^0 \triangleq \{y_0\}.$$

The *information state*, denoted π_t , is the conditional density of state x_t given data ζ^t .

$$\pi_t \triangleq \text{pdf}(x_t | \zeta^t). \quad (3)$$

We further impose the following standing assumption on the random variables and control inputs.

Assumption 1. The signals in (1-2) satisfy:

1. w_t and v_t are i.i.d. sequences with known densities.
2. x_0, w_t, v_t are mutually independent for all $t, l \geq 0$.
3. The control input u_t at time instant $t \geq 0$ is a function of the data ζ^t and given initial state density $\pi_{0|-1}$.

Denoting by $\mathbb{E}_t[\cdot]$ and $\mathbb{P}_t[\cdot]$ the conditional expected value and probability with respect to state x_t – with conditional density π_t – and random variables $\{(w_k, v_{k+1}) : k \geq t\}$, respectively, our goal is to solve the *finite-horizon stochastic optimal control problem* (FHOCP)

$$\mathcal{P}_N(\pi_t) : \begin{cases} \inf_{u_t, \dots, u_{t+N-1}} \mathbb{E}_t \left[\sum_{k=t}^{t+N-1} c(x_k, u_k) + c_N(x_{t+N}) \right], \\ \text{s.t.} \quad x_{k+1} = f(x_k, u_k, w_k), \\ x_t \sim \pi_t, \\ \mathbb{P}_{k+1}[x_{k+1} \in \mathbb{X}_{k+1}] \geq 1 - \epsilon_{k+1}, \\ u_k \in \mathbb{U}_k, \\ k = t, \dots, t+N-1. \end{cases}$$

In theory, solving the FHOCP at each time t and subsequently implementing the first control in a receding horizon fashion leads to a number of desirable closed-loop properties, as discussed in Sehr and Bitmead (2016a). However, solving the FHOCP is computationally intractable in practice, a fact that has led to a number of approaches in MPC for nonlinear stochastic dynamics. We propose a novel strategy that is oriented at the structure of SMPC based on the FHOCP, but numerically tractable at least for low state dimensions.

As a result of the Markovian state equation (1) and measurement equation (2), the optimal control inputs in the FHOCP must inherently be *separated* feedback policies (e.g. Bertsekas (1995); Kumar and Varaiya (1986)). That is, control input u_t depends on the available data ζ^t and initial density $\pi_{0|-1}$ solely through the current information state, π_t . Optimality thus requires propagating π_t and policies g_t , where

$$u_t = g_t(\pi_t). \quad (4)$$

Motivated by this two-component separated structure of stochastic optimal output-feedback control, we propose an extension of the SCMPC approach to nonlinear systems, merged with a numerical approximation of the information state update via particle filtering. Before proceeding with this novel approach, we briefly revisit the two components of SMPC via solution of the FHOCP.

3. STOCHASTIC MODEL PREDICTIVE CONTROL

The information state is propagated via the *Bayesian Filter* (see e.g. Chen (2003); Simon (2006)):

$$\pi_t = \frac{\text{pdf}(y_t | x_t) \pi_{t|t-1}}{\int \text{pdf}(y_t | x_t) \pi_{t|t-1} dx_t}, \quad (5)$$

$$\pi_{t+1|t} \triangleq \int \text{pdf}(x_{t+1} | x_t, u_t) \pi_t dx_t, \quad (6)$$

for $t \in \{0, 1, 2, \dots\}$ and initial density $\pi_{0|-1}$. The recursion (5-6) has the following features:

- The *measurement update* (5) combines the *a priori* conditional density, $\pi_{t|t-1}$, and $p(y_t | x_t)$, derived from (2) using knowledge of: the function $h(\cdot, \cdot)$, the density of v_t , and the value of y_t .
- The *time update* (6) combines π_t and $p(x_{t+1}|x_t, u_t)$, derived from (1) using knowledge of: control input u_t , function $f(\cdot, \cdot, \cdot)$, and the density of w_t .
- For linear Gaussian systems, the filter recursion (5-6) reduces to the well-known Kalman Filter.

Combined with solution of the FHOCP, this leads to the following SMPC algorithm, as discussed in Sehr and Bitmead (2016a).

Algorithm 1 Stochastic Model Predictive Control

- 1: **Offline:**
 - 2: Solve $\mathcal{P}_N(\cdot)$ for the optimal policies, $g_0^*(\cdot), \dots, g_{N-1}^*(\cdot)$.
 - 3: **Online:**
 - 4: **for** $t = 0, 1, 2, \dots$ **do**
 - 5: Measure y_t
 - 6: Compute π_t
 - 7: Apply first optimal control policy, $u_t = g_0^*(\pi_t)$
 - 8: Compute $\pi_{t+1|t}$
 - 9: **end for**
-

Notice how this algorithm differs from common practice in stochastic model predictive control in that it explicitly uses the information states π_t . Throughout the literature, these information states – conditional densities – are commonly replaced by state estimates. While this makes the problem more tractable, one no longer solves the underlying stochastic optimal control problem. The central divergence however lies in Step 2 of the algorithm, in which the SDPE is presumed solved offline for the optimal feedback policies, $g_t(\pi_t)$, from (4). This is an extraordinarily difficult proposition in many cases but captures the optimality, and hence duality, as a closed-loop feedback control law. The complexity of this step lies not only in computing a vector functional but also in the internal propagation of the information state within the SDPE.

4. TRACTABLE NONLINEAR OUTPUT-FEEDBACK MODEL PREDICTIVE CONTROL

In this section, we motivate a novel approach to output-feedback MPC that maintains the separated structure of SMPC while being numerically tractable for modest problem size.

4.1 Approximate Information State & Particle Filter

The BF (5-6) propagates the information state π_t to implement a necessarily separated stochastic optimal output-feedback control law. While implementing this recursion precisely is possible only in special cases such as linear Gaussian systems, where the densities can be finitely parametrized, the BF can be implemented approximately by means of the Particle Filter, with the approximation improving with the number of particles, as described for instance in Simon (2006). In parallel with the BF, the PF consists of two parts: the forward propagation of the state density, and the resampling of the density using the next measurement.

The following algorithm describes a version of the PF amenable to PMPC in the context of this paper. This is a slightly modified version of the filter design described by Simon (2006).

Algorithm 2 Particle Filter (PF)

- 1: Sample N_p particles, $\{x_{0,p}^-, p = 1, \dots, N_p\}$, from density $\pi_{0|-1}$.
 - 2: **for** $t = 0, 1, 2, \dots$ **do**
 - 3: Measure y_t .
 - 4: Compute the relative likelihood q_p of each particle $x_{t,p}^-$ conditioned on the measurement y_t by evaluating $\text{pdf}(y_t | x_{t,p}^-)$ based on (2) and $\text{pdf}(v_t)$.
 - 5: Normalize $q_p \rightarrow q_p / \sum_{p=1}^{N_p} q_p$.
 - 6: Sample N_p particles, $x_{t,p}^+$, via *resampling* based on the relative likelihoods q_p .
 - 7: Given u_t , propagate $x_{t+1,p}^+ = f(x_{t,p}^+, u_t, w_{t,p})$, where $w_{t,p}$ is generated based on $\text{pdf}(w_t)$.
 - 8: **end for**
-

While a number of variations – such as roughening of the particles and differing *resampling* strategies, including importance sampling – of this basic algorithm may be sensible depending on the system at hand, this basic algorithm suffices in presenting a numerical method of approximating the Bayesian Filter to arbitrary degree of accuracy with increasing number of particles N_p (see e.g. Smith and Gelfand (1992)). For a more detailed discussion on the PF for use in state-estimate feedback control, see Rawlings and Mayne (2009).

4.2 Scenario MPC and Particle Model Predictive Control

The Scenario Approach to MPC (See Calafiore and Fagiano (2013); Grammatico et al. (2016); Lee (2014); Mesbah et al. (2014); Schildbach et al. (2014)) commences from state x_t or state estimate, $\hat{x}_{t|t}$. It propagates, i.e. simulates, an open-loop controlled stochastic system with sampled process noise density $\text{pdf}(w_t)$. These propagated samples are then used to evaluate controls for constraint satisfaction and for open-loop optimality with probabilities tied to the sampled w_t densities. In many regards, this is congruent to repeated forward propagation of the PF via (6) without measurement update (5) and commencing from a singular density at x_t or $\hat{x}_{t|t}$. Particle MPC simply replaces the starting point, $\hat{x}_{t|t}$, by the collection of particles $\{x_{t,p}^+, p = 1, \dots, N_p\}$ distributed as π_t , as illustrated in Figure 1.

Before introducing the PMPC algorithm, we define a sampled, particle version of the FHOCP, with N_s scenarios and N_p available *a posteriori* particles at time t ,




Fig. 1. State density evolution in: Scenario MPC calculations (dots and solid outlines) and, Particle MPC (dashed outlines), for three steps into the future.

$$\tilde{\mathcal{P}}_N(\{x_{t,p}^+, p = 1, \dots, N_p\}) : \begin{cases} \inf_{u_t, \dots, u_{t+N-1}} \sum_{s=0}^{N_s} \left(\sum_{k=t}^{t+N-1} c(x_{k,s}, u_k) + c_N(x_{t+N,s}) \right), \\ \text{s.t.} \quad x_{k+1,s} = f(x_{k,s}, u_k, w_{k,s}), \\ x_{t,s} \in \{x_{t,p}^+, p = 1, \dots, N_p\}, \\ \tilde{\mathbb{P}}_{k+1}[x_{k+1} \in \mathbb{X}_{k+1}] \geq 1 - \epsilon_{k+1}, \\ u_k \in \mathbb{U}_k, \\ s = 1, \dots, N_s, \quad k = t, \dots, t+N-1, \end{cases}$$

where the statement

$$\tilde{\mathbb{P}}_{k+1}[x_{k+1} \in \mathbb{X}_{k+1}] \geq 1 - \epsilon_{k+1}$$

means that $x_{k+1,s} \in \mathbb{X}_{k+1}$ for at least $(1 - \epsilon_{k+1})N_s$ scenarios. Following the approach in Schildbach et al. (2013), one may also choose to replace this constraint by $x_{k+1} \in \mathbb{X}_{k+1}$ and select the number of scenarios N_s according to the desired constraint violation levels ϵ_{k+1} . We are now in position to formulate the PMPC algorithm following the schematic in Figure 1.

Algorithm 3 Particle Model Predictive Control (PMPC)

- 1: Generate N_p *a priori* particles, $x_{0,p}^-$, based on $\pi_{0|-1}$.
 - 2: **for** $t = 0, 1, 2, \dots$ **do**
 - 3: Measure y_t .
 - 4: Compute the relative likelihood q_p of each particle $x_{t,p}^-$ conditioned on the measurement y_t by evaluating $\text{pdf}(y_t | x_{t,p}^-)$ based on (2) and $\text{pdf}(v_t)$.
 - 5: Normalize $q_p \rightarrow q_p / \sum_{p=1}^{N_p} q_p$.
 - 6: Generate N_p *a posteriori* particles, $x_{t,p}^+$, via *resampling* based on the relative likelihoods q_p .
 - 7: Solve $\tilde{\mathcal{P}}_N(\{x_{t,p}^+, p = 1, \dots, N_p\})$ for the optimal scenario control values $u_t^*, \dots, u_{t+N-1}^*$.
 - 8: Given u_t^* , propagate $x_{t+1,p}^- = f(x_{t,p}^+, u_t^*, w_{t,p})$, where $w_{t,p}$ is generated based on $\text{pdf}(w_t)$.
 - 9: **end for**
-

4.3 Computational Demand

Computational tractability of PMPC deteriorates with increasing: number of particles, number of scenarios, system dimensions, control signal grid spacing and MPC horizon. Without providing much greater detail and structure, it is

unclear how to select an appropriate numbers of scenarios or particles in the nonlinear case. Suppose the state and input dimensions are n and m and that the numbers of particles and scenarios are chosen as $N_p = P^n$ and $N_s = S^n$ for positive integers P and S (samples per dimension, in effect), respectively, and that the MPC horizon is N . Further assuming a grid of U^m points in the control space and brute-force evaluation of all possible sequences, the order of growth for PMPC is approximately $\mathcal{O}(P^n + S^n U^{mN})$.

Notice that the computational demand associated with the conditional density approximation in PMPC is additive in terms of the overall computational demand. This indicates that, provided the PF is computationally tractable for given state dimensions, tractability of PMPC is roughly equivalent to tractability of standard state-feedback SCMP. In the example below, we found that scenario optimization tends to be the computational bottleneck at least for low system dimensions. Clearly, this observation holds only when the scenario optimization is performed by explicit enumeration of all feasible sequences over a grid in the control space, which may be avoided for particular problem instances. But the experience also confirms that in the nonlinear case the open- or closed-loop control calculation dominates the computational burden in comparison to state estimation.

5. NUMERICAL EXAMPLE

Consider the scalar, nominally unstable nonlinear system

$$x_{t+1} = 1.5x_t + \text{atan}((x_t - 1)^2)u_t + w_t,$$

$$y_t = x_t^3 - x_t + v_t,$$

where x_0, w_t and v_t are mutually independent random variables for all $t, l \geq 0$ and

$$x_0 \sim \mathcal{U}(1, 2), \quad w_t \sim \mathcal{U}(-2, 2), \quad v_t \sim \mathcal{N}(0, 5),$$

for all $t \geq 0$. We aim to minimize the quadratic cost function


$$J_N(\pi_t, u_t, \dots, u_{t+N-1}) = \mathbb{E}_t \left[\sum_{k=t}^{t+N-1} (100x_k^2 + u_k^2) + 100x_{t+N}^2 \right],$$

while satisfying the constraints

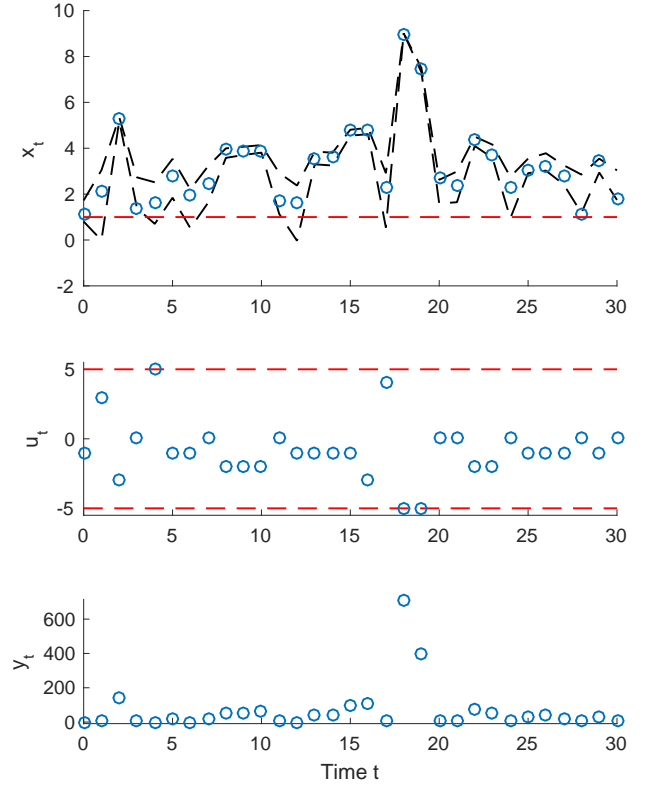
$$\mathbb{P}_{k+1}[x_{k+1} \geq 1] \geq 0.9, \quad -5 \leq u_k \leq 5,$$

along the control horizon N , that is $k \in \{t, \dots, t+N-1\}$ for $t \geq 0$. Notice how this system has both limited observability and controllability close to the constraint but infeasible unconstrained optimal states. In combination with the very noisy measurements, this is a challenging control problem. To implement PMPC as described in Section 4 for this nonlinear stochastic output-feedback control problem, we further restrict the control inputs to integer values, such that $u_t \in \{-5, -4, \dots, 4, 5\}$. Figure 2 displays simulated closed-loop state trajectories, control values and measurement values for four PMPC controllers with differing parameters subject to the same realizations of process and measurement noise, respectively.


Figure 2a displays closed-loop simulation results under PMPC with horizon $N = 3$, $N_p = 5,000$ particles and $N_s = 1,000$ scenarios. While the poor observability properties of the system show close to the probabilistic constraint, it is satisfied at all times in this simulation. This




(a) $N = 3$, $N_p = 5,000$, $N_s = 1,000$.



(b) $N = 3$, $N_p = 100$, $N_s = 1,000$.



(c) $N = 3$, $N_p = 5,000$, $N_s = 50$.



(d) $N = 2$, $N_p = 5,000$, $N_s = 1,000$.

Fig. 2. Simulation data for example in Section 5 over 30 samples, running PMPC with control horizon N , number of particles N_p and number of scenarios N_s . State, control and measurement values (blue), probabilistic and hard constraints (red), 95% confidence interval of PF (black). All controllers are subject to the same realization of the process noise w_k .

is still the case when decreasing the number of particles to $N_p = 100$ in the simulation displayed Figure 2b. However, we see how in this case, the decreased accuracy of the PF leads to larger state-values in closed-loop. Similar behavior is observed in Figure 2c when reducing the number of scenarios to $N_s = 50$. Additionally, the controller violates the probabilistic constraint 3 times in this case. This trend continues when reducing the horizon to $N = 2$, as displayed in Figure 2d.

6. CONCLUSION

We presented PMPC as a novel approach to output-feedback control of stochastic nonlinear systems. Generating scenarios not only from the distribution of the process noise but also from the particles of the PF, PMPC combines the benefits of the PF and SCMPC in a natural fit, allowing for a numerically tractable version of stochastic MPC with general nonlinear dynamics, cost and probabilistic constraints. Given a particular system instance, the algorithm and its properties may be adapted to exploit specific problem structure. Such extensions include: sub-optimal probing via additional constraints, scenario removal algorithms, provable closed-loop properties such as constraint satisfaction with specified confidence levels, and optimization over parametrized control policies.

REFERENCES

- Bertsekas, D.P. (1995). *Dynamic programming and optimal control*. Athena Scientific, Belmont, MA.
- Calafiore, G.C. and Fagiano, L. (2013). Stochastic model predictive control of LPV systems via scenario optimization. *Automatica*, 49(6), 1861–1866.
- Cannon, M., Cheng, Q., Kouvaritakis, B., and Raković, S.V. (2012). Stochastic tube MPC with state estimation. *Automatica*, 48(3), 536–541.
- Chen, Z. (2003). Bayesian filtering: From Kalman filters to particle filters, and beyond. *Statistics*, 182(1), 1–69.
- Chisci, L., Rossiter, J.A., and Zappa, G. (2001). Systems with persistent disturbances: predictive control with restricted constraints. *Automatica*, 37(7), 1019–1028.
- Copp, D.A. and Hespanha, J.P. (2014). Nonlinear output-feedback model predictive control with moving horizon estimation. In *53rd IEEE Conference on Decision and Control*, 3511–3517. Los Angeles, CA.
- Grammatico, S., Zhang, X., Margellos, K., Goulart, P., and Lygeros, J. (2016). A scenario approach for non-convex control design. *IEEE Transactions on Automatic Control*, 61(2), 334–345.
- Kumar, P.R. and Varaiya, P. (1986). *Stochastic Systems: Estimation, Identification, and Adaptive Control*. Prentice-Hall, Englewood Cliffs, NJ.
- Lee, J.H. (2014). From robust model predictive control to stochastic optimal control and approximate dynamic programming: A perspective gained from a personal journey. *Computers & Chemical Engineering*, 70, 114–121.
- Maciejowski, J.M. (2002). *Predictive Control with Constraints*. Prentice Hall, Englewood Cliffs, NJ.
- Marafioti, G., Bitmead, R.R., and Hovd, M. (2014). Persistently exciting model predictive control. *International Journal of Adaptive Control and Signal Processing*, 28(6), 536–552.
- Mayne, D.Q. (2014). Model predictive control: Recent developments and future promise. *Automatica*, 50(12), 2967–2986.
- Mayne, D.Q., Raković, S.V., Findeisen, R., and Allgöwer, F. (2009). Robust output feedback model predictive control of constrained linear systems: Time varying case. *Automatica*, 45(9), 2082–2087.
- Mayne, D.Q., Rawlings, J.B., Rao, C.V., and Scokaert, P.O.M. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, 36(6), 789–814.
- Mesbah, A., Streif, S., Findeisen, R., and Braatz, R.D. (2014). Stochastic nonlinear model predictive control with probabilistic constraints. In *American Control Conference*, 2413–2419. Portland, OR.
- Rawlings, J.B. and Mayne, D.Q. (2009). *Model predictive control: Theory and design*. Nob Hill Publishing, Madison, WI.
- Schildbach, G., Fagiano, L., Frei, C., and Morari, M. (2014). The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations. *Automatica*, 50(12), 3009–3018.
- Schildbach, G., Fagiano, L., and Morari, M. (2013). Randomized solutions to convex programs with multiple chance constraints. *SIAM Journal on Optimization*, 23(4), 2479–2501.
- Schwarm, A.T. and Nikolaou, M. (1999). Chance-constrained model predictive control. *AIChE Journal*, 45(8), 1743–1752.
- Sehr, M.A. and Bitmead, R.R. (2016a). Stochastic model predictive control: Output-feedback, duality and guaranteed performance. Submitted to *Automatica*.
- Sehr, M.A. and Bitmead, R.R. (2016b). Sumptus cohiberi: The cost of constraints in MPC with state estimates. In *American Control Conference*, 901–906. Boston, MA.
- Simon, D. (2006). *Optimal State Estimation: Kalman, H_∞ , and Nonlinear Approaches*. John Wiley & Sons, New York, NY.
- Smith, A.F.M. and Gelfand, A.E. (1992). Bayesian statistics without tears: a sampling-resampling perspective. *The American Statistician*, 46(2), 84–88.
- Sui, D., Feng, L., and Hovd, M. (2008). Robust output feedback model predictive control for linear systems via moving horizon estimation. In *American Control Conference*, 453–458. Seattle, WA.
- Yan, J. and Bitmead, R.R. (2005). Incorporating state estimation into model predictive control and its application to network traffic control. *Automatica*, 41(4), 595–604.